Essential Circles and Gromov-Hausdorff Convergence of Covers

Conrad Plaut
Department of Mathematics
University of Tennessee
Knoxville, TN 37996
cplaut@math.utk.edu

Jay Wilkins
Department of Mathematics
University of Connecticut
196 Auditorium Road, Unit 3009
Storrs, CT 06269-3009
leonard.wilkins@uconn.edu

April 4, 2013

Abstract

We give various applications of essential circles (introduced in [15]) in a compact geodesic space X. Essential circles completely determine the homotopy critical spectrum of X, which we show is precisely $\frac{2}{3}$ the covering spectum of Sormani-Wei. We use finite collections of essential circles to define "circle covers", which generalize the δ-covers of Sormani-Wei (equivalently the ε-covers of [15]). We show that, unlike δ-covers, circle covers are in a sense closed with respect to Gromov-Hausdorff convergence, and we prove a finiteness theorem concerning deck groups of circle covers that does not hold for covering maps in general. This theorem allows us to completely understand the structure of Gromov-Hausdorff limits of δ-covers. Essential circles also give rise to a certain collection of subgroups of the fundamental group called *circle groups*, which give an algebraic enhancement of the covering/homotopy critical spectra. Finally, we use essential circles to strengthen a theorem of E. Cartan by finding a new (even for compact Riemannian manifolds) finite set of generators of the fundamental group of a semilocally simply connected compact geodesic space. We conjecture that there is always a generating set of this sort having minimal cardinality among all generating sets.

Keywords: Gromov-Hausdorff convergence, essential circles, covering maps, fundamental group, geodesic space, length spectrum, Laplace spectrum

1 Introduction

Christina Sormani and Guofang Wei ([16],[17],[18],[19]), and Valera Berestovskii and Conrad Plaut ([1],[2]) studied covering space constructions that encode geometric information. Sormani-Wei utilized a classical construction of Spanier ([20]) that provides a covering map $\pi^{\delta}:\widetilde{X}^{\delta}\rightarrow X$ corresponding to the open cover of a geodesic space X by open δ -balls, which they called the δ-cover of X. Berestovskii-Plaut developed a new method for uniform spaces (hence any metric space) that utilizes discrete chains and homotopies rather than traditional paths and homotopies. The special case of metric spaces was developed further by Plaut and Wilkins in [15]. For any connected metric space X and $\varepsilon>0$, this construction yields a covering map $\phi_{\varepsilon}:X_{\varepsilon}\rightarrow X$, with deck group $\pi_{\varepsilon}(X)$ that is a kind of fundamental group at a given scale. In this paper we show that, despite the very different construction methods, when $\delta=\frac{3\varepsilon}{2}$ and X is a geodesic space, these two covering maps are naturally isometrically equivalent (Corollary 15). Given this equivalence, and to avoid confusion, in this paper we will use the terminology and notation of [15], including when mentioning results of Sormani-Wei.

One characteristic of ε-covers observed by Sormani-Wei is that they are not "closed" with respect to Gromov-Hausdorff convergence of compact geodesic spaces in the following sense: it may happen that a sequence of compact geodesic spaces X_i converges to a compact space X in the Gromov-Hausdorff sense, while the possibly non-compact covers $(X_i)_{\varepsilon}$ do not converge to X_{ε} in the pointed Gromov-Hausdorff sense. Sormani-Wei did show that there is always a subsequence of the $(X_i)_{\varepsilon}$ that pointed converges to a space Y that covers X and is covered by X_{ε} (i.e. the ε-covers of a convergent sequence of compact geodesic spaces are *precompact* - Proposition 7.3, [17]). Until now, however, nothing more has been known about the limiting space Y. The main purpose of this paper is to introduce a new, metrically determined kind of covering space that extends the notion of ε-cover and is closed, in a certain sense, with respect to Gromov-Hausdorff convergence. In fact, we are able to completely understand what happens to not only these covers but also to their corresponding deck groups when the base spaces converge.

The idea, without technical details, is as follows. The fact that ε-covers are not closed is related to what we called in [15] the *homotopy critical values* of a space, which for a compact geodesic space are the same (up to a multiplied factor of $\frac{3}{2}$) as the *covering spectrum* of Sormani-Wei (Corollary 15). For compact geodesic spaces, the homotopy critical values are discrete in $(0,\infty)$ (Theorem 11, [15]) and indicate precisely when the equivalence type of the ε-covers changes as ε varies. For example, given a standard geodesic (i.e. Riemannian) circle of length $a>0$, the ε-covers are all isometries when $\varepsilon>\frac{a}{3}$, but are the standard universal cover when $\varepsilon\leq\frac{a}{3}$, i.e. the homotopy critical spectrum of this circle is $\{\frac{a}{3}\}$. In general one may imagine ε as sliding from the diameter of the space towards 0, as the space "unrolls" a bit more at each homotopy critical value. If the space is semilocally simply connected, then the process stabilizes with the universal cover, and the smallest homotopy critical value is $\frac{1}{3}$ of the 1-systole

of the space (smallest non-null homotopic closed geodesic - Corollary 43, [15]). If the space is not semilocally simply connected, then the "unrolling" process may never end, although by using an inverse limit one will obtain the *uniform universal cover* defined in [2].

Now suppose X_i is a geodesic circle of length $a - \frac{1}{i}$, so $X_i \to X$, where X is the circle of length a. But $(X_i)_{\frac{a}{3}} = X$, while $X_{\frac{a}{3}} = \mathbb{R}$. Of course in this case $(X_i)_{\frac{a}{3}}$ *does* converge to some X_ε, in fact for any $\varepsilon > \frac{a}{3}$. To see how it is possible for $(X_i)_\varepsilon$ to not converge to any ε-cover at all, one needs to involve multiplicity of homotopy critical values, most simply illustrated as follows: Suppose that Y denotes the flat torus obtained by identifying the sides of a rectangle of dimensions $0 < 3a \leq 3b$. When $a < b$, a and b are distinct homotopy critical values: For $\varepsilon > b$, $Y_\varepsilon = Y$, for $a < \varepsilon \leq b$, Y_ε is a flat metric cylinder over a circle of length $3a$, and for $\varepsilon \leq a$, Y_ε is the plane. When $a = b$, the torus unrolls immediately into the plane at $\varepsilon = a$; the ε-covers "skip" the cylinder. There is a natural notion of *multiplicity* of a homotopy critical value ([15]), and in this case a is a homotopy critical value of multiplicity 2, since both topological holes are of the same size and are detected by the ε-covers simultaneously.

Now take a sequence of tori T_i obtained from $(1 - \frac{1}{i}) \times (1 + \frac{1}{i})$-rectangles. So T_i has homotopy critical values $\frac{1}{3} - \frac{1}{3i}$ and $\frac{1}{3} + \frac{1}{3i}$ with single multiplicity, while the limiting torus has a critical value $\frac{1}{3}$ with multiplicity 2. In fact, $(T_i)_{\frac{1}{3}}$ is a cylinder for all i , with limit Y a cylinder of circumference 1, which as observed above is not an ε-cover of the limiting torus. What happens here (and Theorem 1 shows it is true is true in general), is that distinct homotopy critical values merge in the limit to a single homotopy critical value with multiplicity. The covers that we describe next allow us to tease apart the multiplicity to find the "missing" intermediate covers.

In [15] we showed that $\varepsilon > 0$ is a homotopy critical value of a compact geodesic space X if and only if X has what we called an *essential ε-circle*, which is a special kind of closed geodesic that we will describe in more detail later. Multiplicity of homotopy critical values is defined using a natural equivalence between essential circles; i.e. the multiplicity of ε is the number of equivalence classes of essential ε-circles.

Now suppose that $\varepsilon > 0$ and $\mathcal{T}$ is any finite collection of essential δ-circles such that $\delta \geq \varepsilon$. We define a natural normal subgroup $K_\varepsilon(\mathcal{T})$ of $\pi_\varepsilon(X)$ (which is the trivial group when $\mathcal{T} = \varnothing$) that acts freely and properly discontinuously on X_ε with quotient $X_\varepsilon^{\mathcal{T}}$. Then there is a natural induced mapping $\phi_\varepsilon^{\mathcal{T}} : X_\varepsilon^{\mathcal{T}} \to X$ which is also a covering map with deck group naturally isomorphic to $\pi_\varepsilon^{\mathcal{T}}(X) = \pi_\varepsilon(X)/K_\varepsilon(\mathcal{T})$. We call $\phi_\varepsilon^{\mathcal{T}}$ the $(\mathcal{T}, \varepsilon)$-cover of X . In general we will refer to these covers as *circle covers*. Note that when $\mathcal{T} = \varnothing$, $\phi_\varepsilon^{\mathcal{T}} = \phi_\varepsilon$, so circle covers extend the notion of ε -covers. Note also that $\phi_\varepsilon^{\mathcal{T}}$ in general depends on ε as well as $\mathcal{T}$. In the next statements, the arrow "$\to$" indicates Gromov-Hausdorff convergence (respectively, pointed Gromov-Hausdorff convergence) of the compact spaces X_i (respectively, for the possibly non-compact covering spaces). The first theorem below is a special (but essential and strengthened) case used to prove the second.

Theorem 1 *Suppose that $X_i \to X$, where each X_i is compact geodesic and let $\varepsilon > 0$. Then for any $\delta < \varepsilon$ sufficiently close to ε, $(X_i)_\delta \to (X)_\varepsilon$ and $\pi_\delta(X_i)$ is isomorphic to $\pi_\varepsilon(X)$ for all large i. In particular, if ε is not a homotopy critical value of X then $(X_i)_\varepsilon \to X_\varepsilon$ and $\pi_\varepsilon(X_i)$ is eventually isomorphic to $\pi_\varepsilon(X)$.*

Theorem 2 *Suppose that $X_i \to X$, where each X_i is compact geodesic, and for each i there are an $\varepsilon_i > 0$ and a finite collection $\mathcal{T}_i$ of essential τ-circles in X_i such that $\tau \geq \varepsilon_i$. If $\{\varepsilon_i\}$ has a positive lower bound then for any positive $\varepsilon \leq \liminf \varepsilon_i$ there exist a subsequence $\{X_{i_k}\}$ and a finite collection $\mathcal{T}$ of essential τ-circles in X with $\tau \geq \varepsilon$, such that $(X_{i_k})_{\varepsilon_{i_k}}^{\mathcal{T}_{i_k}} \to X_\varepsilon^{\mathcal{T}}$ and $\pi_{\varepsilon_{i_k}}^{\mathcal{T}_{i_k}}(X_{i_k})$ is isomorphic to $\pi_\varepsilon^{\mathcal{T}}(X)$ for all large k.*

Theorem 2 shows that circle covers are a natural extension of ε-covers that contain all Gromov-Hausdorff limits of ε-covers. The situation is more complicated without a positive lower bound on the size of the essential circles. In general, there may be no convergent subsequence of covers. For example, if H denotes the geodesic Hawaiian Earring, then as $\varepsilon \to 0$ the ε-covers H_ε of H contain graphs of higher and higher valency (see [2] for more details). Even if the covers do converge, the limiting cover may not be a circle cover. For instance, consider the torus $T_i = S^1_{3/i} \times S^1_1$ formed by circles of circumference $\frac{3}{i}$ and 1. Then $T_i \to S^1_1$. If we choose ε_i so that $\varepsilon_i < \frac{1}{i}$ and $\varepsilon_i \to 0$, then each $(T_i)_{\varepsilon_i}$ for $i \geq 3$ is the universal cover $\mathbb{R}^2$, but of course $\mathbb{R}^2$ is not any kind of cover of S^1. This particular example satisfies the assumptions of Theorem 1.1, [9], in which Ennis and Wei showed the following: Suppose $X_i \to X$ are all compact geodesic spaces having (categorical, possibly not simply connected) universal covers. The latter assumption is equivalent to the homotopy critical spectra having positive lower bounds $\varepsilon_i, \varepsilon$, and in fact the universal covers are $(X_i)_{\varepsilon_i}$ and X_ε, respectively. Theorem 1.1, [9] says that if one additionally assumes that the spaces X_i have dimension uniformly bounded above and the spaces X_{ε_i} pointed Gromov-Hausdorff converge to a space $\overline{X}$, then there is a subgroup $H \subset Iso(\bar{X})$ such that $\bar{X}/H$ is the universal cover of X. One interesting aspect of this result is that, as the previously mentioned example of collapsing tori shows, the subgroup H need not be discrete.

Our proofs use the discrete methods developed in [2] and [15], which are much more amenable to this kind of geometric problem than classical continuous methods. For example, Gromov-Hausdorff convergence of compact metric spaces is characterized by the existence of almost isometries (or σ-isometries) that generally are not continuous, and therefore classical methods using continuous paths are not easy to apply. However, discrete methods allow one to easily see that for any function $f : X \to Y$ between metric spaces, an induced function $f_\# : X_\varepsilon \to Y_\delta$ exists provided f only satisfies a kind of continuity at a single scale: for any $x, y \in X$, if $d(x,y) < \varepsilon$ then $d(f(x), f(y)) < \delta$ (see also Theorem 2 in [2]). Moreover, we show that when f happens to be a σ-isometry, $f_\#$ is in fact a quasi-isometry (see Remark 26 and Theorem 30) with distortion constants explicitly controlled by σ. These induced quasi-isometries allow us to

gain control over the Gromov-Hausdorff distance between balls in the (possibly non-compact) spaces X_ε and Y_δ in explicit terms, and prove our main theorems.

The discrete analog of the essential ε-circle is the essential ε-triad. An ε-triad is characterized (in compact geodesic spaces) by the fact that any triple of equidistant points on an essential circle is an essential triad, and if one joins each pair of points on an essential triad by a geodesic, then the resulting triangle is an essential circle. Therefore these notions are equivalent, but essential triads are much easier to work with in the setting of Gromov-Hausdorff convergence than essential circles. In fact, in our more precise definitions below, we use essential triads to define essential circles and therefore the collections $\mathcal{T}$ described above.

As a byproduct of this work we use essential circles to create a new (even for Riemannian manifolds) set of generators of the fundamental group of a compact geodesic space, which we conjecture can always be taken to be minimal (Theorem 23). Some questions also naturally arise. For example, is it possible to characterize circle covers among all covers (this extends a question from [15] about characterizing ε-covers)? We know that not all covers of a compact geodesic space are circle covers. For example, it follows from our results that the geodesic circle discussed above has only two circle covers: the trivial cover and the universal cover. Therefore, other non-equivalent covers of the circle, such as double covers, cannot be circle covers. But at this point we are only able to identify when a cover is a circle cover in the following ways: (1) by exclusion when we know all the circle covers in a particular example, (2) if the cover is explicitly defined as a circle cover, or (3) if it is known to be so by Theorem 2. In this connection we note that the natural analog of Theorem 2 for covering maps in general is not true. In fact, for a circle cover π of a compact geodesic space X, a lower bound $\varepsilon > 0$ on the size of the circles is equivalent to being covered by the ε-cover of X. Now let $\psi_k : C_k \to C_1$ be the k-fold cover of the geodesic circle C_1 by the geodesic circle of length k, which as we have mentioned is not a circle cover for $k > 1$. Each of these covers is covered by the universal covering space of C_1, which is the $\frac{1}{3}$-cover of C_1. It is also true that $C_k \to \mathbb{R}$ in the pointed Gromov-Hausdorff sense (in fact it is not hard to show in general precompactness of covering spaces covered by an ε-cover). However, the deck groups of these covering maps are $\mathbb{Z}_k$, which of course are all distinct.

Another question of interest is related to the fairly old question concerning the degree to which various spectra (Laplacian, length, covering) determine geometric properties in a compact geodesic space, including whether they must be isometric. Note that, up to a multiplied constant, the covering and *homotopy critical* spectra (i.e. the collection of homotopy critical values) are contained in the length spectrum. While this was already observed by Sormani-Wei, this is an immediate consequence of the previously mentioned fact that X contains an essential ε-circle if and only if ε is a homotopy critical value. The relationship between the length and the Laplace spectra was first considered in [3], [5], [8]. Already de Smit, Gornet, and Sutton have shown that the covering spectrum is not a spectral invariant ([6], [7]) by extending Sunada's method [21] to determine when two manifolds have the same Laplace spectrum. However, essential circles allow one to enhance the notion of covering/homotopy critical spectrum in the

following way. Given a compact geodesic space X, each circle covering of X corresponds to a subgroup of $\pi_1(X)$, which we will call a *circle group*. Specifically, a circle group is the kernel of the natural map $\Lambda : \pi_1(X) \to \pi_\varepsilon(X)$ mentioned at the beginning of the next section, composed with the quotient map from $\pi_\varepsilon(X)$ to $\pi_\varepsilon(X)/K_\varepsilon(\mathcal{T})$ described above. The collection of all circle groups, partially ordered by inclusion, provides an algebraic refinement of the homotopy critical spectrum which in principle should say more about how similar two spaces are. That is, what can be said about compact geodesic spaces that not only share the same homotopy critical spectra, including multiplicity, but also share the same partially ordered collection of circle groups up to isomorphism?

As one final question, we ask whether the partially ordered set of circle groups is a topological invariant. We know already that the subgroups of the fundamental group corresponding to the covering maps ϕ_ε can depend on the metric, for example as a result of multiplicity in the case of tori described above. But we introduced circle groups specifically to handle multiplicity, and at present we do not have examples of different geodesic metrics on the same compact topological space that have different circle groups.

2 Background and Examples

We begin with some background from [15]; proofs and further details regarding the results in this section may be found there. In a metric space X, and for $\varepsilon > 0$, an ε-chain is a finite sequence $\{x_0, ..., x_n\}$ such that for all i, $d(x_i, x_{i+1}) < \varepsilon$. An ε-homotopy consists of a finite sequence $\langle \gamma_0, ..., \gamma_n \rangle$ of ε-chains, where each γ_i differs from its predecessor by a "basic move": adding or removing a *single* point, always leaving the endpoints fixed. Fixing a basepoint $*$, X_ε is defined to be the set of all ε-homotopy equivalence classes of ε-chains starting at $*$, and $\phi_\varepsilon : X_\varepsilon \to X$ is the endpoint map. In a connected space, choice of basepoint is immaterial, so we will not include it in our notation and will assume that all maps are base-point preserving.

The group $\pi_\varepsilon(X)$ is the subset of X_ε consisting of classes of ε-loops starting and ending at $*$ with operation induced by concatenation, i.e., $[\alpha]_\varepsilon * [\beta]_\varepsilon = [\alpha * \beta]_\varepsilon$. We denote the reversal of a chain α by $\overline{\alpha}$. As expected, for $[\alpha]_\varepsilon \in \pi_\varepsilon(X)$, $([\alpha]_\varepsilon)^{-1} = [\overline{\alpha}]_\varepsilon$, and the identity is $[*]_\varepsilon$.

For any ε-chain $\alpha = \{x_0, ..., x_n\}$, we set $\nu(\alpha) := n$ and define its *length* by

$$L(\alpha) := \sum_{i=1}^{n} d(x_i, x_{i-1}).$$

Defining $|[\alpha]_\varepsilon| := \inf\{L(\gamma) : \gamma \in [\alpha]_\varepsilon\}$ leads to a metric on X_ε given by

$$d([\alpha]_\varepsilon, [\beta]_\varepsilon) := |[\overline{\alpha} * \beta]_\varepsilon| = \inf\{L(\kappa) : \alpha * \kappa * \overline{\beta} \text{ is } \varepsilon\text{-null homotopic}\} \qquad (1)$$

This metric has a number of nice properties that we will need. For example, $\pi_\varepsilon(X)$ acts on X_ε as isometries via the map induced by preconcatenation by an ε-loop. If one defines $\phi_\varepsilon : X_\varepsilon \to X$ to be the endpoint mapping (which

is well-defined since ε-homotopies preserve endpoints), then this function is a local isometry and, provided X is connected, a regular covering map with deck group naturally identified with $\pi_\varepsilon(X)$. When X happens to be a geodesic space (which will soon be our underlying assumption) then so is X_ε, and in fact the above metric coincides with the traditional lifted geodesic metric on the covering space X_ε (Proposition 23, [15]). The definition using (1) is very useful for our purposes, but since we will need the lifted geodesic metric for arbitrary covering spaces, we will recall the definition now. Given a covering map $\phi : X \to Y$, where Y is a geodesic space and X is connected, the lifted geodesic metric on X is defined by $d(x, y) = \inf\{L(c) : \phi \circ c$ such that c is a path joining x and $y\}$. As pointed out in [15], geodesic metrics are uniquely determined by their local values, and in particular the lifted geodesic metric is uniquely determined by the fact that ϕ is a local isometry.

There is a mapping from fixed-endpoint homotopy classes of continuous paths to ε-homotopy classes of ε-chains defined as follows: For any continuous path $c : [0,1] \to X$, choose $0 = t_0 < \cdots < t_n = 1$ fine enough that every image $c([t_i, t_{i+1}])$ is contained in the open ball $B(c(t_i), \varepsilon)$. Then the chain $\{c(t_0), ..., c(t_n)\}$ is called a subdivision ε-chain of c. Setting $\Lambda([c]) := [c(t_0), ..., c(t_n)]_\varepsilon$ produces a well-defined function that is length non-increasing in the sense that $|\Lambda([c])| \leq |[c]| := \inf\{L(d) : d \in [c]\}$. Restricting Λ to the fundamental group at any base point yields a homomorphism $\pi_1(X) \to \pi_\varepsilon(X)$ that we will also refer to as Λ. When X is geodesic, Λ is surjective since the successive points of an ε-loop λ may be joined by geodesics to obtain a path loop whose class goes to $[\alpha]_\varepsilon$. The kernel of Λ is precisely described by Corollary 16. Variations of Λ and their applications to the fundamental group and universal covers are further examined by the second author in [22].

A partial inverse operation to Λ is given by the following notion: Let $\alpha := \{x_0, ..., x_n\}$ be an ε-chain in a metric space X, where $\varepsilon > 0$. A *stringing* of α consists of a path $\widehat{\alpha}$ formed by concatenating paths γ_i from x_i to x_{i+1} where each path γ_i lies entirely in $B(x_i, \varepsilon)$. If each γ_i is a geodesic then we call $\widehat{\alpha}$ a *chording* of α. Note that by "geodesic" in this paper we mean an arclength-parameterized path whose length is equal to the distance between its endpoints, and not a locally minimizing path as is the more common meaning in Riemannian geometry. We will need the following two basic results.

Proposition 3 *If α is an ε-chain in a chain connected metric space X then the unique lift of any stringing $\widehat{\alpha}$ starting at the basepoint $[*]_\varepsilon$ in X_ε has $[\alpha]_\varepsilon$ as its endpoint.*

Corollary 4 *If α and β are ε-chains in a chain connected metric space X such that there exist stringings $\widehat{\alpha}$ and $\widehat{\beta}$ that are path homotopic then α and β are ε-homotopic.*

We also need some basic technical results. The first of these quantifies the idea that "uniformly close" ε -chains are ε-homotopic. Of course "close" depends on ε. Given $\alpha = \{x_0, ..., x_n\}$ and $\beta = \{y_0, ..., y_n\}$ with $x_i, y_i \in X$,

define $\Delta(\alpha, \beta) := \max_i \{d(x_i, y_i)\}$. For any $\varepsilon > 0$, if α is an ε-chain we define $E_\varepsilon(\alpha) := \min_i \{\varepsilon - d(x_i, x_{i+1})\} > 0$. When no confusion will result we will eliminate the ε subscript.

Proposition 5 *Let X be a metric space and $\varepsilon > 0$. If $\alpha = \{x_0, ..., x_n\}$ is an ε-chain and $\beta = \{x_0 = y_0, ..., y_n = x_n\}$ is such that $\Delta(\alpha, \beta) < \frac{E(\alpha)}{2}$ then β is an ε-chain that is ε-homotopic to α.*

The reader will likely have noticed that the previous proposition requires that the chains in question have the same number of points. The next lemma shows that this is not really an issue. It is useful in many ways–for example to find "convergent subsequences" of classes of chains, much like a discrete version of Ascoli's Theorem.

Lemma 6 *Let $L, \varepsilon > 0$ and α be an ε-chain in a metric space X with $L(\alpha) \leq L$. Then there is some $\alpha' \in [\alpha]_\varepsilon$ such that $L(\alpha') \leq L(\alpha)$ and $\nu(\alpha') = \left\lfloor \frac{2L}{\varepsilon} + 1 \right\rfloor$.*

For any $\delta \geq \varepsilon > 0$, every ε-chain (respectively ε-homotopy) is also a δ-chain (respectively δ-homotopy) and there is a well-defined mapping $\phi_{\delta\varepsilon} : X_\varepsilon \to X_\delta$ given by $\phi_{\delta\varepsilon}([\alpha]_\varepsilon) = [\alpha]_\delta$. When X is geodesic, this mapping is also a regular covering map and local isometry, though for general metric spaces it may not be surjective. Restricting the map $\phi_{\delta\varepsilon}$ to the group $\pi_\delta(X)$ induces a homomorphism $\theta_{\delta\varepsilon} : \pi_\varepsilon(X) \to \pi_\delta(X)$, which is injective if and only if $\phi_{\delta\varepsilon}$ is.

A number $\varepsilon > 0$ is called a *homotopy critical value* for X if there is an ε-loop α based at $*$ such that α is not *ε-null* (i.e. ε-homotopic to the trivial chain) but is δ-null for all $\delta > \varepsilon$. We have the following essential connection between homotopy critical values and ε-covers:

Lemma 7 *If X is a geodesic space then the covering map $\phi_{\varepsilon\delta} : X_\delta \to X_\varepsilon$ is injective if and only if there are no homotopy critical values σ with $\delta \leq \sigma < \varepsilon$.*

Corollary 8 *If λ is an ε-loop in a geodesic space X of length less than 3ε then λ is ε-null.*

If $\alpha = \{x_0, .., x_n\}$ is a chain in a geodesic space X then a refinement of α consists of a chain β formed by inserting between each x_i and x_{i+1} some subdivision chain of a geodesic joining x_i and x_{i+1}. If β is an ε-chain we will call β an ε-refinement of α. Note that if α is an ε-chain, then any ε-refinement of α is ε-homotopic to α, and, hence, any two ε-refinements of α are ε-homotopic. A special case is a *midpoint refinement*, which simply uses a midpoint between pairs of points. Of course refinements always exist in geodesic spaces, but not in general metric spaces. Refinements are in important tool when working with convergence questions, since the property of being an ε-chain is an "open condition" and may not be preserved when passing from a sequence of chains to a limit.

Definition 9 *If X is a metric space and $\varepsilon > 0$, an ε-loop of the form $\lambda = \alpha * \tau * \overline{\alpha}$, where $\nu(\tau) = 3$, will be called ε-small. Note that this notation includes the case when α consists of a single point–i.e. $\lambda = \tau$.*

The next proposition was stated in [15] with the assumption that $\varepsilon < \delta$, but the same proof works for $\varepsilon = \delta$, and we will need that case in the present paper. In essence it "translates" a homotopy into a product of small loops.

Proposition 10 *Let X be a geodesic space and $0 < \varepsilon \leq \delta$. Suppose α, β are ε-chains and $\langle \alpha = \gamma_0, ..., \gamma_n = \beta \rangle$ is a δ -homotopy. Then $[\beta]_\varepsilon = [\lambda_1 * \cdots * \lambda_r * \alpha * \lambda_{r+1} * \cdots * \lambda_n]_\varepsilon$, where each λ_i is an ε -refinement of a δ-small loop.*

An ε-*triad* in a geodesic space X is a triple $T := \{x_0, x_1, x_2\}$ such that $d(x_i, x_j) = \varepsilon$ for all $i \neq j$; when ε is not specified we will simply refer to a triad. We denote by α_T the loop $\{x_0, x_1, x_2, x_0\}$. We say that T is *essential* if some (equivalently any) ε -refinement of α_T is not ε-null. The equivalence of "some" and "any" in the preceding definition - as well as the useful fact that the ε-refinement of α_T may always be taken to be a midpoint refinement - follow from Proposition 37 in [15]. Essential ε-triads T_1 and T_2 are defined to be *equivalent* if some (equivalently any) ε-refinement of α_{T_1} is freely ε-homotopic to an ε-refinement of either α_{T_1} or $\overline{\alpha_{T_1}}$, and, again, it suffices to consider only midpoint refinements. See [15] for the definition of "free ε-homotopy", which is analogous to the classical meaning for paths.

As mentioned in the Introduction, *essential circles* may be defined as geodesic triangles obtained by joining the points of an essential triad by geodesics. An essential circle determined by an essential ε-triad is referred to as an essential ε-circle. Equivalence of the underlying triads is used to define equivalence of essential circles. Note that equivalent essential circles may not be freely path homotopic due to "small holes" that block traditional homotopies but not ε-homotopies. Essential circles are necessarily the images of non-null, closed geodesics that are shortest in their respective homotopy classes, but they have an even stronger property: an essential circle is metrically embedded in the sense that its metric as a subspace of X is the same as the intrinsic metric of the circle (Theorem 39, [15]).

3 Covering Maps and Quotients

We begin by recalling some results concerning quotients of metric spaces. In this paper, all actions are by isometries and are discrete in the sense of [14]. Discreteness of an action is a uniform version of properly discontinuous, which is implied by the following property for any G acting by isometries on a metric space Y : There exists some $\varepsilon > 0$ such that for all $y \in Y$ and non-trivial $g \in G$, $d(y, g(y)) \geq \varepsilon$. Note that in the case where Y is the ε-cover of a metric space X, $\pi_\varepsilon(X)$ acts discretely on X_ε, since the restriction of ϕ_ε to any $B([\alpha]_\varepsilon, \frac{\varepsilon}{2}) \subset X_\varepsilon$ is an isometry onto its image in X ([15]). When Y is geodesic, then the quotient metric on Y/G (cf. [13]) is the uniquely determined geodesic metric such that

the quotient mapping is a local isometry ([15]). Combining this observation with Proposition 28 in [14] we obtain the following:

Proposition 11 *Suppose that X is a geodesic space, G acts discretely by isometries on X, and H is a normal subgroup of G. Then G/H acts discretely by isometries on X/H via $gH(Hx) = g(x)H$ and the mapping $Gx \mapsto (G/H)(Hx)$ an isometry from X/G to $(X/H)/(G/H)$.*

Among basic applications we have that for any geodesic space X and $0 < \delta < \varepsilon$, the covering map $\phi_\varepsilon : X_\varepsilon \to X$ is isometrically equivalent to the induced mapping $\zeta : X_\delta / \ker \theta_{\delta\varepsilon} \to X$. Here $\ker \theta_{\delta\varepsilon}$ acts discretely and isometrically as a normal subgroup of $\pi_\delta(X)$ with $X = X_\delta/\pi_\delta(X)$, and ζ is the unique covering map such that $\zeta \circ \pi = \phi_\delta$, while $\pi : X_\delta \to X_\delta / \ker \theta_{\delta\varepsilon}$ is the quotient mapping.

We know already from the results of [2] that if X is a compact metric space, Y is connected, and $f : Y \to X$ is a covering map then for small enough $\varepsilon > 0$ there is a covering map $g : X_\varepsilon \to Y$. The next proposition refines this statement when X is geodesic.

Proposition 12 *Let X be a compact geodesic space and suppose that $f : Y \to X$ is a covering map, where Y is connected. Suppose that $\varepsilon > 0$ is at most $\frac{2}{3}$ of a Lebesgue number for a covering of X by open sets evenly covered by f. Then there is a covering map $g : X_\varepsilon \to Y$ such that $\phi_\varepsilon = f \circ g$.*

Proof. Choose a basepoint $*$ in Y such that $f(*) = *$ and define $g([\alpha]_\varepsilon)$ to be the endpoint lift of some stringing $\widehat{\alpha}$ starting at $*$ in Y. We need to check that g is well-defined. By iteration, it suffices to prove the following: If $\alpha := \{x_0, ..., x_n\}$ and α' is an ε-chain $\{x_0, ..., x_i, x, x_{i+1}, ..., x_n\}$, then for any stringings $\widehat{\alpha}$ and $\widehat{\alpha'}$, the endpoints of the lifts of $\widehat{\alpha}$ and $\widehat{\alpha'}$ starting at $*$ are the same. Let $\{\gamma_i\}$ and $\{\gamma_i'\}$ be geodesics joining x_i, x_{i+1} and β_1, β_2 be geodesics from x_i to x and x to x_{i+1}, respectively. Note that each of the loops formed by each γ_i and the reversal of γ_i' has diameter smaller than ε and thus lifts as a loop in Y. Moreover, the triangle formed by the geodesics β_1, β_2, and either γ_i or γ_i' has diameter less than $\frac{3\varepsilon}{2}$ and hence lifts as a loop. This proves that g is well-defined. It now follows from a standard result in topology that g is a covering map ([12]). (This result is stated with the additional assumption that all maps are continuous, but that assumption is superfluous because the third mapping - g in this case - is locally a homeomorphism.) ■

The next theorem shows that any covering space of a compact geodesic space can be obtained in a particularly natural way as a quotient of the space X_ε obtained in Proposition 12.

Theorem 13 *Let X be a compact geodesic space and suppose that $f : Y \to X$ is a covering map, where Y is connected and has the lifted geodesic metric from X. Let $\varepsilon > 0$ be such that there is a covering map $g : X_\varepsilon \to Y$ with $\phi_\varepsilon = f \circ g$. Define a subgroup K of $\pi_\varepsilon(X)$ by*

$$K := \{[\lambda]_\varepsilon : \exists\ \kappa \in [\lambda]_\varepsilon \text{ such that some stringing } \widehat{\kappa} \text{ of } \kappa \text{ lifts as a loop in } Y\}.$$

Then

1. $K = \{[\lambda]_\varepsilon : \textit{for all } \kappa \in [\lambda]_\varepsilon, \textit{ every stringing } \widehat{\kappa} \textit{ of } \kappa \textit{ lifts as a loop to } Y\}$.
2. *There is a covering equivalence* $\phi : Y \to X_\varepsilon/K$ *such that* $\pi = \phi \circ g$, *where* $\pi : X_\varepsilon \to X_\varepsilon/K$ *is the quotient map.*
3. K *is a normal subgroup if and only if* f *is a regular covering map.*

Proof. The first part follows from Proposition 3. Define $\phi(y) := \pi(x)$, where $x \in g^{-1}(y)$. To see why ϕ is well-defined, suppose that $g([\alpha]_\varepsilon) = g([\beta]_\varepsilon) = y$. Since the lifts $\widetilde{\alpha}$ and $\widetilde{\beta}$ of $\widehat{\alpha}$ and $\widehat{\beta}$ to the basepoint in X_ε have $[\alpha]_\varepsilon$ and $[\beta]_\varepsilon$ as their endpoints (Proposition 3), $g \circ \widetilde{\alpha}$ and $g \circ \widetilde{\beta}$ both end in y. But these two curves are the unique lifts of $\widehat{\alpha}$ and $\widehat{\beta}$ to Y, so they both end in y. In other words if $\lambda := \alpha * \overline{\beta}$, then $[\lambda]_\varepsilon \in K$. Moreover, $[\alpha]_\varepsilon := [\lambda]_\varepsilon * [\beta]_\varepsilon$, so $[\alpha]_\varepsilon$ and $[\beta]_\varepsilon$ lie in the same orbit of K, and $\pi([\alpha]_\varepsilon) = \pi([\beta]_\varepsilon)$. This shows that ϕ is well-defined, and clearly $\pi = \phi \circ g$.

For surjectivity, let $K[\alpha]_\varepsilon \in X_\varepsilon/K$ and define $y := g([\alpha]_\varepsilon)$. We have

$$\phi(y) = \phi(g([\alpha]_\varepsilon)) = \pi([\alpha]_\varepsilon) = K[\alpha]_\varepsilon.$$

For injectivity, suppose that $y, z \in Y$ satisfy $\phi(y) = \phi(z)$. Then for $[\alpha]_\varepsilon, [\beta]_\varepsilon$ such that $g([\alpha]_\varepsilon) = y$ and $g([\beta]_\varepsilon) = z$ we have that $[\alpha]_\varepsilon, [\beta]_\varepsilon$ lie in the same orbit, so $[\alpha]_\varepsilon * [\overline{\beta}]_\varepsilon = [\lambda]_\varepsilon \in K$. Consequently, the lifts of stringings $\widehat{\alpha}$ and $\widehat{\beta}$ to Y at the basepoint must have the same endpoint. But one lift ends in y and the other ends in z, so $y = z$. This shows that ϕ is a homeomorphism. The natural covering map $\xi : X_\varepsilon/K \to X$ is defined by $\xi(K[\alpha]_\varepsilon) = \phi_\varepsilon([\alpha]_\varepsilon)$ and therefore $\xi \circ \phi = f$, showing that ϕ is a covering equivalence.

If K is normal then by Lemma 39 in [14], the covering map $\xi : X_\varepsilon/K \to X$ is a topological quotient map with respect to a well-defined induced action of $\pi_\varepsilon(X)/K$, defined by $[\lambda]_\varepsilon K([\alpha]_\varepsilon K) = ([\lambda]_\varepsilon([\alpha])_\varepsilon)K$. Since ξ is also a covering map, it follows from standard results in topology that ξ is regular. Since ϕ is a covering equivalence, f is also regular. Conversely, suppose that ξ is regular and let H denote its group of covering transformations. Define a function $h : \pi_\varepsilon(X) \to H$ as follows: Given $[\alpha]_\varepsilon \in \pi_\varepsilon(X)$, let y be the endpoint of the lift of some chording $\widehat{\alpha}$ to Y starting at the basepoint. Since $\phi_\varepsilon = g \circ f$, by uniqueness y must also be the endpoint $f \circ c$, where c is the lift of $\widehat{\alpha}$ to X_ε . By Proposition 3, the endpoint of c is just $[\alpha]_\varepsilon$, and hence $y = f([\alpha]_\varepsilon)$ depends only on $[\alpha]_\varepsilon$. That is, if we let $h([\alpha]_\varepsilon)$ be the (unique) $\mu \in H$ such that $\mu(*) = y$, then h is well-defined. Sorting through the definition shows that h is a homomorphism with kernel K, since y is the basepoint exactly when $\widehat{\alpha}$ lifts as a loop. ■

Corollary 14 *Let* X *be compact geodesic and* $\varepsilon > 0$. *Then an* ε*-loop* α *is* ε*-null if and only if some (equivalently any) stringing of* α *lifts to a loop in* X_ε.

Proof. We only note that α is ε-null if and only if $\beta * \alpha * \overline{\beta}$ is ε-null for any ε-chain from the basepoint to the starting point of α, and the analogous statement also holds for null-homotopies of any stringing of α. ■

The δ-covering map defined by Sormani-Wei is obtained using a construction of Spanier that provides for any open cover $\mathcal{U}$ of a connected, locally path connected topological space Z a covering map $\phi_{\mathcal{U}} : W \to Z$. The covering map is characterized by the fact that a path loop c at the basepoint in Z lifts as a loop to W if and only if its homotopy equivalence class $[c]$ lies in the subgroup $S_{\mathcal{U}}$ of $\pi_1(Z)$, which we will call the *Spanier Group of* $\mathcal{U}$, generated by all loops of the following form: $c * L * \overline{c}$, where c is a path loop starting at the basepoint and L is a path loop lying entirely in some set in the open cover $\mathcal{U}$. For their construction, Sormani-Wei took $\mathcal{U}$ to be the cover of X by open δ-balls. We will denote the corresponding Spanier Group by S_δ.

Corollary 15 *For any geodesic space* X *and* $\delta = \frac{3\varepsilon}{2} > 0$ *, there is an equivalence of the covering maps (hence an isometry)* $h : \widetilde{X}^\delta \to X_\varepsilon$.

Proof. An immediate consequence of the definition of $\widetilde{X}^\delta$ is that all open δ-balls are evenly covered by π^δ, and hence we may take δ for the Lebesgue number in Proposition 12. That proposition gives a covering map $\psi : X_\varepsilon \to \widetilde{X}^\delta$ and Theorem 13 will finish the proof if we can show that the group K for $Y := \widetilde{X}^\delta$ is trivial. If $[\lambda]_\varepsilon \in K$ then by definition some chording $\widehat{\lambda}$ lifts as a loop in $\widetilde{X}^\delta$. In other words, $\widehat{\lambda}$ is homotopic to a concatenation of paths of the form $c_i * L_i * \overline{c_i}$ where L_i is a path loop that lies entirely in an open δ-ball. According to Corollary 4 we need only show that any refinement ε -chain of any such path is ε-null. In other words, it is enough to show the following: Any ε-refinement of a rectifiable loop f in an open ball $B(x, \delta)$ is ε-null. Given such a loop f, the distance from points on f to x has a maximum $D < \delta$. Now subdivide f into segments σ_i whose endpoints x_i, x_{i+1} satisfy $d(x_i, x_{i+1}) < \delta - D$. Then each of the ε-loops $\kappa_i := \{x, x_i, x_{i+1}, x\}$ (where $x_n = x_0$ for the highest index n) has length less than $D + D + \delta - D = D + \delta < 2\delta = 3\varepsilon$. But then each κ_i, and hence f, is ε-null by Corollary 8. ∎

Corollary 16 *If* X *is a geodesic space,* $\varepsilon > 0$*, and* $\Lambda : \pi_1(X) \to \pi_\varepsilon(X)$ *is the natural homomorphism defined in the second section, then* $\ker \Lambda = S_{\frac{3\varepsilon}{2}}$.

Proof. We have $[c] \in \ker \Lambda$ if and only if some subdivision ε -chain α of c is ε-null. Equivalently, by Corollary 14, any stringing of α lifts to a loop in X_ε. But since $X_\varepsilon = \widetilde{X}^\delta$ by Corollary 15, this is equivalent to $[c] \in S_{\frac{3\varepsilon}{2}}$. ∎

Definition 17 *Suppose* $\mathcal{T}$ *is a finite collection of essential triads in a compact geodesic space and* $\varepsilon > 0$ *is such that each* $T \in \mathcal{T}$ *is a* δ*-triad for some* $\delta \geq \varepsilon$*. Define* $K_\varepsilon(\mathcal{T})$ *to be the subgroup of* $\pi_\varepsilon(X)$ *generated by the collection* $\Gamma_\varepsilon(\mathcal{T})$ *of all* $[\alpha * T' * \overline{\alpha}]_\varepsilon$*, where* α *is an* ε*-chain starting at* $*$ *and* T' *is an* ε*-refinement of* $T \in \mathcal{T}$*. Finally, we define* $\phi_\varepsilon^{\mathcal{T}} : X_\varepsilon / K_\varepsilon(\mathcal{T}) \to X$ *by* $\phi_\varepsilon^{\mathcal{T}}(K_\varepsilon(\mathcal{T})[\alpha]_\varepsilon) := \phi_\varepsilon([\alpha]_\varepsilon)$*. We will call a covering map equivalent to some* $\phi_\varepsilon^{\mathcal{T}}$ *a circle covering map (including the case when* $\mathcal{T}$ *is empty, in which case we take* $K_\varepsilon(\mathcal{T})$ *to be the trivial group, so* $\phi_\varepsilon^{\mathcal{T}} = \phi_\varepsilon$*).*

Remark 18 *First, note that $K_\varepsilon(\mathcal{T})$ is normal in $\pi_\varepsilon(X)$, since any conjugate of $[\alpha * T' * \alpha^{-1}]_\varepsilon$ has the same form. If $\mathcal{T} = \{T_i\}_{i=1}^n$, it is also easy to check that for any fixed choice of ε-chains α_i from $*$ to T_i, $K_\varepsilon(\mathcal{T})$ is the smallest normal subgroup containing the finite set $\{[\alpha_i * T_i' * \overline{\alpha_i}]_\varepsilon\}_{i=1}^n$. However, we do not know in general whether $K_\varepsilon(\mathcal{T})$ is finitely generated. Finally, a word of caution. While δ-refinements of an essential δ -triad T are all δ-homotopic, different ε-refinements of T need not be ε-homotopic when δ is larger than ε. In fact, as simple geodesic graphs show, the vertices of T may be joined by different geodesics that together form loops that are always δ-null but are not ε-null. That is, in general, replacing a δ-triad $T \in \mathcal{T}$ with a δ-equivalent δ-triad may change the group $K_\varepsilon(\mathcal{T})$. This further emphasizes the essential dependence of $K_\varepsilon(\mathcal{T})$ on not just the collection $\mathcal{T}$ but on the value of ε.*

Notation 19 *When convenient we will denote X by X_∞; this makes sense since any chain can be considered as an ∞-chain and every sequence of chains in which one point is removed or added to get from one chain to the next is an ∞-homotopy. Then every chain is ∞-homotopy equivalent to the chain $\{x, y\}$ where x and y are its endpoints, and hence the mapping $\phi_\varepsilon : X_\varepsilon \to X$ is naturally identified with $\phi_{\infty\varepsilon} : X_\varepsilon \to X_\infty$. This saves us from having to consider the mapping ϕ_ε as a special case in the statements that follow.*

Proposition 20 *Let $\varepsilon > 0$ be a homotopy critical value for a compact geodesic space X. Then there is some $\delta > \varepsilon$ such that if $\{x_0, x_1, x_2, x_0\}$ is δ-small with a midpoint refinement α that is not ε-null then α is ε -homotopic to a midpoint refinement of an essential ε-triad.*

Proof. If the statement were not true then there would exist $\left(\varepsilon + \frac{1}{i}\right)$-small loops $\{x_i, y_i, z_i, x_i\}$ having midpoint subdivision chains $\mu_i = \{x_i, m_i, y_i, n_i, z_i, p_i, x_i\}$ that are not ε-null but are not ε-homotopic to a midpoint refinement of an essential ε-triad. By taking subsequences if necessary, we may suppose that $\{x_i, m_i, y_i, n_i, z_i, p_i, x_i\} \to \{x, m, y, n, z, p, x\}$, where $\{x, y, z, x\}$ is an ε -small loop which is not ε-null. Hence $\{x, y, z\}$ must be an essential ε-triad. Moreover, for large i, $\{x_i, m_i, y_i, n_i, z_i, p_i, x_i\}$ is ε-homotopic to a midpoint subdivision of $\{x, y, z, x\}$, a contradiction. ∎

Definition 21 *In a geodesic space X, we will call a path (resp. ε-chain) of the form $k * c * \overline{k}$, where k is a path (resp. ε-chain), $\overline{k}$ is its reversal, and c is an arclength parameterization of an essential circle (resp. a midpoint refinement of an essential ε-triad), a lollipop (resp. ε-lollichain). If the path k in the lollipop is locally length minimizing (possibly not minimal) then we call the lollipop a geodesic lollipop.*

Note that for fixed ε, the homomorphism $\Lambda : \pi_1(X) \to \pi_\varepsilon(X)$ maps classes of lollipops determined by essential ε-circles to ε-lollichains. In fact, if c is an essential ε-circle determined by an essential ε-triad T then we can choose a midpoint refinement β of α_T such that each point of β lies on c. By definition, $[c]$ is mapped via Λ to $[\beta]_\varepsilon$. Conversely, if β is a midpoint refinement of an essential

ε-triad T, we can define an essential ε-circle c containing β by joining the points of β by geodesics. Then Λ will map $[c]$ to $[\beta]_\varepsilon$. Anchoring the essential circles to the base point by adjoining paths k and choosing corresponding ε-chains κ so that $[\kappa]_\varepsilon = \Lambda([k])$ yields the full conclusion.

Theorem 22 *Let X be a compact geodesic space, $0 < \varepsilon < \delta \leq \infty$, and $\mathcal{T}$ any (possibly empty!) collection that contains a representative for every essential τ-triad with $\varepsilon \leq \tau < \delta$. Then $\ker \theta_{\delta\varepsilon} = K_\varepsilon(\mathcal{T})$. Consequently, $\pi_\delta(X) = \pi_\varepsilon(X)/K_\varepsilon(\mathcal{T})$, and the covering map $\phi_{\delta\varepsilon} : X_\varepsilon \to X_\delta$ is equivalent to the quotient covering map $\pi : X_\varepsilon \to X_\varepsilon/K_\varepsilon(\mathcal{T})$.*

Proof. First of all, note that the inequality $\tau < \delta$ shows that each element of $K_\varepsilon(\mathcal{T})$ is δ-null and hence $K_\varepsilon(\mathcal{T}) \subset \ker \theta_{\delta\varepsilon}$. For the opposite inclusion, let $[\lambda]_\varepsilon \in \ker \theta_{\delta\varepsilon}$, meaning that λ is δ-null. We will start with the case when ε is a homotopy critical value of X and $\delta > \varepsilon$ is close enough to ε that $\delta < 2\varepsilon$ and Proposition 20 is valid: whenever $\{x_0, x_1, x_2, x_0\}$ is δ-small with a midpoint refinement α that is not ε-null then α is ε-homotopic to a midpoint refinement of an essential ε-triad. By Proposition 10 λ is ε-homotopic to a product of midpoint refinements λ_i of δ-small loops. Since Proposition 20 holds, each λ_i is either ε-null or ε-homotopic to a non-null ε-lollichain. That is, $[\lambda]_\varepsilon \in K_\varepsilon(\mathcal{T})$.

Next, observe that if there are no homotopy critical values τ with $\varepsilon \leq \tau < \delta$, then on the one hand $\mathcal{T}$ must be empty, and on the other hand, $\theta_{\delta\varepsilon}$ is an isomorphism so its kernel is trivial, and we are finished. Suppose now that there is a single critical value τ between ε and δ , which, by the previous case, may be assumed to satisfy $\varepsilon \leq \tau < \delta$. We may choose δ_1 with $\tau < \delta_1 < \delta$ satisfying the requirement of the special case proved in the first paragraph to obtain that $\ker \theta_{\delta_1\tau} = K_\tau(\mathcal{T})$. Now both $\theta_{\tau\varepsilon}$ and $\theta_{\delta\delta_1}$ are isomorphisms, and $\theta_{\delta\varepsilon} = \theta_{\delta\delta_1} \circ \theta_{\delta_1\tau} \circ \theta_{\tau\varepsilon}$ by definition. Therefore,

$$\ker \theta_{\delta\varepsilon} = \theta_{\tau\varepsilon}^{-1}(\ker \theta_{\delta_1\tau}) = \theta_{\tau\varepsilon}^{-1}(K_\tau(\mathcal{T})) = K_\varepsilon(\mathcal{T}).$$

For the general case we have $\varepsilon \leq \varepsilon_i < \cdots < \varepsilon_j < \delta := \varepsilon_{j+1}$ where $\{\varepsilon_i, ..., \varepsilon_j\}$ is the set of all homotopy critical values between ε and δ, which has at least two elements. For $i \leq k \leq j$, let $\mathcal{T}_k$ be the set of all ε_k -triads in $\mathcal{T}$, so that $\mathcal{T} = \cup \mathcal{T}_k$. By the previous case we have for all k that $\ker \theta_{\varepsilon_{k+1}\varepsilon_k} = K_{\varepsilon_k}(\mathcal{T}_k)$. If $[\lambda]_\varepsilon \in \ker \theta_{\delta\varepsilon} = \ker \theta_{\varepsilon_{j+1}\varepsilon}$ then $x := \theta_{\varepsilon_j\varepsilon}([\lambda]_\varepsilon) \in \ker \theta_{\varepsilon_{j+1}\varepsilon_j} = K_{\varepsilon_j}(\mathcal{T}_j)$, so we may write x as a finite product $\Pi_r[\beta_r^j]_{\varepsilon_j}$ with where each β_r^j is an ε_j-lollichain made from an ε_j-refinement of some element of $\mathcal{T}_j$. Now let λ_r^j be any ε-refinement of β_r^j. By definition, $[\lambda_r^j]_\varepsilon \in K(\mathcal{T})$ for all r, and $\theta_{\varepsilon_j\varepsilon}([\lambda_r^j]_\varepsilon) = [\beta_r^j]_{\varepsilon_j}$. Therefore we may write $[\lambda]_\varepsilon = [\lambda'_j]_\varepsilon \left(\Pi_r[\lambda_r^j]_\varepsilon\right)$ for some $[\lambda'_j]_\varepsilon \in \ker \theta_{\varepsilon\varepsilon_j}$. Since $[\lambda'_j]_\varepsilon \in \ker \theta_{\varepsilon\varepsilon_j}$ we may repeat the same argument to write $[\lambda'_j]_\varepsilon$ as a product of some $[\lambda''_j]_\varepsilon \in \ker \theta_{\varepsilon\varepsilon_{j-1}}$ and a finite product of elements of $K(\mathcal{T})$ (consisting of ε-lollichains formed using ε -refinements of elements of $\mathcal{T}_{j-1}$). After finitely many iterations of this argument we obtain that $[\lambda]_\varepsilon \in K(\mathcal{T})$. ■

It is a straightforward extension of a theorem of E. Cartan that in a compact semi-locally simply connected geodesic space, every path contains a shortest

path in its fixed-endpoint homotopy class, and that path is a locally minimizing geodesic. Likewise, every path loop contains a shortest element in its free homotopy class, and this curve is a closed geodesic. (This is not true in general without semilocal simple connectivity.) Consequently, in such spaces the fundamental group is generated by homotopy classes of loops of the form $\alpha * c * \bar{\alpha}$, where c is a closed geodesic that is shortest in its homotopy class, and α is a locally minimizing geodesic. Now in the case $\delta = \infty$ in Theorem 22, since we know from [15] that $\pi_\varepsilon(X)$ is finitely generated, we obtain that $\pi_\varepsilon(X)$ is generated by finitely many ε-lollichains. If X is semilocally simply connected then for ε small enough, $\pi_\varepsilon(X)$ is isomorphic to $\pi_1(X)$ via the map $\Lambda : \pi_1(X) \to \pi_\varepsilon(X)$ (c.f. [15] or [22]). By the discussion following Definition 21, applying Λ^{-1} to a basis of ε-lollichains give a basis of lollipops for $\pi_1(X)$, and we may replace any lollipop by a geodesic lollipop, up to homotopy equivalence. We thus have shown:

Theorem 23 *Let X be a compact geodesic space. Then $\pi_\varepsilon(X)$ is either trivial or is generated by a finite collection $[\lambda_1]_\varepsilon, ..., [\lambda_n]_\varepsilon$, where each λ_i is an ε-refinement of a δ-lollichain with $\delta \geq \varepsilon$. In particular, if X is semilocally simply connected and not simply connected then $\pi_1(X)$ is generated by a finite collection of equivalence classes of geodesic lollipops.*

Recalling that the notion of essential circle is stronger than just being a non-null, closed geodesic that is shortest in its homotopy class, we see that Theorem 23 is stronger than Cartan's result even in the Riemannian case. In fact, Example 44 of [15] shows that such closed geodesics need not be essential circles even when they are shortest in their homotopy class in a Riemannian manifold.

We conjecture that there is always a collection of δ-lollichains that gives a generating set of $\pi_\varepsilon(X)$ having minimal cardinality.

Proposition 24 *Let X be a compact geodesic space, $0 < \delta < \varepsilon$, and $\mathcal{T}$ be a collection of essential τ-triads such that $\tau \geq \varepsilon$ for each element of $\mathcal{T}$. Let $\mathcal{S} = \mathcal{T} \cup \mathcal{T}'$, where $\mathcal{T}'$ consists of one representative of each essential τ-triad with $\delta \leq \tau < \varepsilon$. Then*

1. *The covering map $\phi_\varepsilon^{\mathcal{T}} : X_\varepsilon^{\mathcal{T}} \to X$ is isometrically equivalent to $\phi_\delta^{\mathcal{S}} : X_\delta^{\mathcal{S}} \to X$ and*

2. *$\pi_\varepsilon^{\mathcal{T}}(X)$ is isomorphic to $\pi_\delta^{\mathcal{S}}(X)$.*

Proof. We will use Proposition 11. First observe that $K_\delta(\mathcal{T}')$, as a normal subgroup of $\pi_\delta(X)$, is a normal subgroup of $K_\delta(\mathcal{S})$. By Theorem 22, $\ker \theta_{\varepsilon\delta} = K_\delta(\mathcal{T}')$ and therefore we may identify the action of $\pi_\varepsilon(X)$ on X_ε with the action of $K_\delta(\mathcal{T})/K_\delta(\mathcal{T}')$ on $X_\delta/K_\delta(\mathcal{T}') = X_\varepsilon$. In order to apply Proposition 11 and finish the proof of the first part, we need to show that $\theta_{\varepsilon\delta}(K_\delta(\mathcal{S})) = K_\varepsilon(\mathcal{T})$. Since $\mathcal{T} \subset \mathcal{S}$, $K_\varepsilon(\mathcal{T}) = \theta_{\varepsilon\delta}(K_\delta(\mathcal{T})) \subset \theta_{\varepsilon\delta}(K_\delta(\mathcal{S}))$. On the other hand, let $[\lambda]_\varepsilon \in \theta_{\varepsilon\delta}(K_\delta(\mathcal{S}))$. By definition, $[\lambda]_\varepsilon = [\lambda_1]_\varepsilon * \cdots * [\lambda_k]_\varepsilon$, where each $\lambda_i = \alpha_i * \beta_i * \overline{\alpha_i}$, α_i is a δ -chain and each β_i is either in $\mathcal{T}$ or $\mathcal{T}'$. But if $\beta_i \in \mathcal{T}'$ then β_i is ε-null,

so $[\lambda_i]_\varepsilon$ is trivial and we may therefore eliminate $[\lambda_i]_\varepsilon$ from the product. The remaining terms are all in $K_\varepsilon(\mathcal{T})$.

For the second part, note that we have shown both $\ker\theta_{\varepsilon\delta} = K_\delta(\mathcal{T}') \subset K_\delta(\mathcal{S})$ and $\theta_{\varepsilon\delta}(K_\delta(\mathcal{S})) = K_\varepsilon(\mathcal{T})$. Therefore, from a basic theorem in algebra we may conclude that $\pi_\varepsilon^{\mathcal{T}}(X) = \pi_\varepsilon(X)/K_\varepsilon(\mathcal{T})$ is isomorphic to $\pi_\delta(X)/K_\delta(\mathcal{S}) = \pi_\delta^{\mathcal{S}}(X)$. ■

4 Gromov-Hausdorff Convergence

Definition 25 *Suppose $f : X \to Y$ is a function between metric spaces and σ is a first degree polynomial with non-negative coefficients. We say that f is a σ-isometry if for all $x, y \in X$, $z \in Y$,*

1. $|d(x,y) - d(f(x), f(y))| \leq \sigma(d(x,y))$ *and*
2. $d(z, f(w)) \leq \sigma(0)$ *for some $w \in X$.*

We refer to the first condition as "distortion at most σ".

If $\sigma = 0$ then a σ-isometry is an isometry, and if $\sigma = \varepsilon > 0$ is constant then Definition 25 agrees with the notion of an ε-isometry given in [4]. If X and Y are compact of diameter at most R, then any σ-isometry is a $\sigma(R)$-isometry. In fact, if X and Y are compact, σ is constant, and d denotes the Gromov-Hausdorff distance, then $d(X,Y) < 2\sigma$ if there is a σ-isometry $f : X \to Y$, and such a σ-isometry exists if $d(X,Y) < \frac{\sigma}{2}$ (Corollary 7.3.28, [4]). In other words, for purposes involving convergence of compact spaces we might as well use constant functions σ. However, our extended definition is needed to study the induced mapping $f_\# : X_\delta \to Y_\varepsilon$ since X_σ and Y_ε are not generally compact even when X and Y are.

Remark 26 *Recall that a quasi-isometry (c.f. [4]) is a map $f : X \to Y$ such that $f(X)$ is a D-net in Y for some $D > 0$ (i.e. for every $y \in Y$ we have $d(y, f(x)) < D$ for some $x \in X$) and $\frac{1}{\lambda}d(x,y) - C \leq d(f(x), f(y)) \leq \lambda d(x,y) + C$ for all x, y and some constants $\lambda \geq 1$, $C > 0$. If $\sigma(x) = mx + b$ then a σ-isometry is a quasi-isometry with $\lambda := \frac{1+m}{1-m}$ and $C = b$, with $\lambda \to 1$ as $m \to 0$. However, it is simpler for our purposes to use Definition 25.*

Proposition 27 *Suppose that (X_i, x_i), (X, x) are proper geodesic spaces and $f_i : X_i \to X$ is a basepoint preserving σ_i -isometry for all i, where σ_i is a first degree polynomial with $\sigma_i \to 0$ pointwise. Then (X_i, x_i) is pointed Gromov-Hausdorff convergent to (X, x).*

Proof. It suffices to show that for any constants $0 < \sigma < 1 < R$, there is a σ-isometry $g_i : B(x_i, R) \to B(x, R)$ for all large i. For large i, $\sigma_i(4R) < \frac{\sigma}{4}$, and the restriction of f_i to $B(x_i, 2R)$ has distortion less than $\frac{\sigma}{4}$. In particular, if $y \in B(x_i, R)$ then $f_i(y) \in B(x, R + \frac{\sigma}{4})$. If $f_i(y) \in B(x, R)$, let $g_i(y) := f_i(y)$. Otherwise, since X is geodesic, there is some $u \in B(x, R)$ such that $d(f_i(y), u) <$

$\frac{\sigma}{4}$ (i.e. u is on a geodesic from $f_i(y)$ to x); in this case define $g_i(y) := u$. By the triangle inequality, the distortion of g_i on $B(x_i, R)$ is at most $\frac{3\sigma}{4}$.

To finish, we need only check condition 2 of the definition of σ-isometry for the constant σ. If $\sigma_i(t) = m_i t + b_i$, then by our choice of i (and $R > 1$), we have $m_i, b_i < \frac{\sigma}{4} < \frac{1}{4}$. Let $z \in B(x, R)$; since X is geodesic we may find $z' \in B(x, R - \frac{\sigma}{2})$ such that $d(z, z') < \frac{\sigma}{2}$. Since $\sigma_i(0) = b_i < \frac{\sigma}{4}$, there is some $w \in X_i$ such that $d(f_i(w), z') < \frac{\sigma}{4}$. By the triangle inequality, $f_i(w) \in B(x, R - \frac{\sigma}{4})$, and hence by definition of g_i, $g_i(w) = f_i(w)$. Again by the triangle inequality, $d(g_i(w), z) \leq \frac{3\sigma}{4} < \sigma$, and we are left only to show that $w \in B(x_i, R)$. Set $d(w, x_i) := S$. Then

$$S \leq d(f_i(w), x) + \sigma_i(S) < R - \frac{\sigma}{4} + \sigma_i(S) < R - \frac{\sigma}{4} + \frac{S}{4} + \frac{\sigma}{4},$$

which implies $S < \frac{4}{3}R$. Since now $w \in B(x_i, 2R)$, we can use the distortion of σ_i to on this ball to improve this estimate and complete the proof:

$$S < d(f_i(w), f_i(x_i)) + \frac{\sigma}{4} = d(f_i(w), x) + \frac{\sigma}{4} < R - \frac{\sigma}{4} + \frac{\sigma}{4} = R.$$

■

Let $f : X \to Y$ be a function between metric spaces. We will extend the notion of the induced function $f_\#$ from [2] as follows. For any metric space Z and $\mu > 0$, let $\overline{Z_\mu}$ consist of all $[\alpha]_\mu$, where α is a μ-chain (not necessarily starting at $*$). For any chain $\alpha = \{x_0, \ldots, x_n\}$ in X, we let $f(\alpha) := \{f(x_0), \ldots, f(x_n)\}$. Note that $f(\alpha * \beta) = f(\alpha) * f(\beta)$ whenever the first concatenation is defined. Suppose now that for some fixed $\varepsilon, \delta > 0$ and all $x, y \in X$, if $d(x, y) < \varepsilon$ then $d(f(x), f(y)) < \delta$. If α is an ε-chain in X then $f(\alpha)$ is a δ-chain in Y. Moreover, if $\eta = \{\eta_1, \ldots, \eta_n\}$ is an ε -homotopy in X then $f(\eta) := \{f(\eta_1), \ldots, f(\eta_n)\}$ is a δ-homotopy in Y. It follows that the mapping $f_\# : \overline{X_\varepsilon} \to \overline{Y_\delta}$ defined by $f_\#([\alpha]_\varepsilon) = [f(\alpha)]_\delta$ is well-defined and satisfies

$$f_\#([\alpha * \beta]_\varepsilon) = f_\#([\alpha]_\varepsilon) * f_\#([\beta]_\varepsilon) \tag{2}$$

whenever the first concatenation is defined.

The next technical proposition sorts through the ways in which $f_\#$ inherits the properties of a σ-isometry f. The statement is not the most general possible - for example it is possible to consider non-constant σ and the first part doesn't depend on the full distortion assumption - but the statement is complicated enough as it is and we will not need more general statements for this paper.

Proposition 28 *Let X, Y be metric spaces, $0 < \frac{\varepsilon}{2} < \omega < \delta < \varepsilon$. Suppose that $f : X \to Y$ is basepoint preserving σ-isometry for some constant σ with $0 \leq \sigma < \min\{\varepsilon - \delta, \frac{\delta - \omega}{4}\}$. Then the induced map $f_\# : \overline{X_\delta} \to \overline{Y_\varepsilon}$ is defined, and the following hold.*

1. *For any δ-chain α in X,*

$$|f_\#([\alpha]_\delta)| \leq |[\alpha]_\delta| + \sigma\left(\frac{4\,|[\alpha]_\delta|}{\varepsilon} + 1\right).$$

2. *For any ω-chain α in X,*

$$|[f(\alpha)]_{\omega+\sigma}| \geq |[\alpha]_\delta| - \sigma\left(\frac{4}{\varepsilon} + \frac{16\sigma}{\varepsilon^2}\right)(|[\alpha]_\omega|) - \sigma\left(\frac{4\sigma}{\varepsilon} + 1\right).$$

3. *If β is an $(\omega - 3\sigma)$-chain in Y starting at $*$ then there exists some ω-chain α in X starting at $*$ such $d(f_\#([\alpha]_\delta), [\beta]_\varepsilon) < \sigma$ in Y_ε.*

Proof. That the induced map is defined follows from $\delta + \sigma < \varepsilon$. According to Lemma 6, up to δ-homotopy and without increasing the length of α, we may assume that $\alpha = \{x_0, \ldots, x_n\}$ where $n = \left\lfloor \frac{2L(\alpha)}{\delta} + 1 \right\rfloor$. We have

$$L(f(\alpha)) \leq L(\alpha) + n\sigma = L(\alpha) + \left(\frac{2L(\alpha)}{\delta} + 1\right)\sigma.$$

By definition, $f(\alpha) \in f_\#([\alpha]_\delta)$, therefore $|f_\#([\alpha]_\delta)| \leq L(f(\alpha))$. The proof of part 1 follows by taking the infimum of the right side and using $\delta > \frac{\varepsilon}{2}$.

For the second part, fix $\tau > 0$ and let $\alpha' := \{x_0, \ldots, x_m\}$ be an ω-chain such that $[\alpha']_\omega = [\alpha]_\omega$ and $L(\alpha') \leq |[\alpha]_\omega| + \tau$. By Lemma 6 we may suppose that $m = \left\lfloor \frac{2L(\alpha')}{\omega} + 1 \right\rfloor$. Then $\beta := f(\alpha') = \{f(x_0), \ldots, f(x_m)\}$ is an $(\omega + \sigma)$ -chain of length at most $K = L(\alpha') + m\sigma$. Let $\eta = \langle \beta = \eta_0, \ldots, \eta_n = \beta' \rangle$ be an $(\omega + \sigma)$-homotopy, where $L(\beta') \leq L(\beta)$; again by Lemma 6 we may assume that $\nu(\beta') = m' := \left\lfloor \frac{2K}{\omega+\sigma} + 1 \right\rfloor$. Denote by y_{ij} the i^{th} point of η_j; note that the points y_{ij} are not necessarily distinct! Nonetheless, we may iteratively choose for each y_{ij} a point $x_{ij} \in X$ such that $d(f(x_{ij}), y_{ij}) < \sigma$ and the following are true: $x_{i0} = x_i$ (which is possible since $\eta_0 = \beta = f(\alpha)$), and if $y_{ij} = y_{ab}$ then $x_{ij} = x_{ab}$. For any x_{ij}, x_{ab},

$$d(x_{ij}, x_{ab}) \leq d(f(x_{ij}), f(x_{ab})) + \sigma \leq d(y_{ij}, y_{ab}) + 3\sigma$$

and therefore if we let η'_j denote the chain in X having x_{ij} as its i^{th} point, $\eta' := \langle \eta'_0, \ldots, \eta'_n \rangle$ is an $(\omega + 4\sigma)$-homotopy between α' and a chain α'' of length at most $L(\beta') + m'\sigma$. Since $\omega + 4\sigma < \delta$, η' is in fact a δ -homotopy. That is,

$$|[\alpha]_\delta| \leq L(\beta') + \sigma\left(\frac{2K}{\omega + \sigma} + 1\right) < L(\beta') + \sigma\left(\frac{4K}{\varepsilon} + 1\right)$$

(in the second inequality we used $\omega + \sigma > \omega > \frac{\varepsilon}{2}$). Next,

$$K \leq L(\alpha') + \sigma\left(\frac{4L(\alpha')}{\varepsilon} + 1\right) = L(\alpha')\left(\frac{4\sigma}{\varepsilon} + 1\right) + \sigma.$$

Putting these two together yields:

$$\begin{aligned}|[\alpha]_\delta| &\leq L(\beta') + \sigma\left(\frac{4}{\varepsilon} + \frac{16\sigma}{\varepsilon^2}\right)L(\alpha') + \sigma\left(\frac{4\sigma}{\varepsilon} + 1)\right) \\ &\leq L(\beta') + \sigma\left(\frac{4}{\varepsilon} + \frac{16\sigma}{\varepsilon^2}\right)(|[\alpha]_\omega| + \tau) + \sigma\left(\frac{4\sigma}{\varepsilon} + 1)\right)\end{aligned}$$

Letting $\tau \to 0$ gives us

$$L(\beta') \geq |[\alpha]_\delta| - \sigma\left(\frac{4}{\varepsilon} + \frac{16\sigma}{\varepsilon^2}\right)(|[\alpha]_\omega|) - \sigma\left(\frac{4\sigma}{\varepsilon} + 1)\right).$$

Taking the infimum over all β' in $[f(\alpha)]_{\omega+\sigma}$ yields the second inequality.

For the third part, let $\beta = \{* = y_0, \ldots, y_n\}$. As usual we take $n := \left\lfloor \frac{2L(\beta)}{\omega} + 1 \right\rfloor$. Since f is a σ-isometry, we may find points $z_i = f(x_i)$ such that $d(z_i, y_i) \leq \sigma$ for $1 \leq i \leq n$. We have

$$d(x_i, x_{i+1}) \leq d(z_i, z_{i+1}) + \sigma \leq d(y_i, y_{i+1}) + \sigma + 2\sigma < \omega$$

and therefore the chain $\alpha := \{* = x_0, \ldots, x_n\}$ is an ω-chain such that $f(\alpha) = \beta' := \{* = z_0, \ldots, z_n\}$. Moreover,

$$L(\alpha) \leq L(\beta) + 3\sigma\left(\frac{2L(\beta)}{\omega} + 1\right) \leq L(\beta) + 3\sigma\left(\frac{4L(\beta)}{\varepsilon} + 1\right).$$

Next, let $\beta'' := \{* = z_0, ..., z_{n-1}, z_n, y_n\}$ and $\beta''' := \{* = y_0, \ldots, y_n, y_n\}$; note that $[\beta''']_\varepsilon = [\beta]_\varepsilon$. Since β''' is an $(\omega - 3\sigma)$-chain, $E(\beta''') > \varepsilon - (\omega - 3\sigma) > 2\sigma$, and since $\Delta(\beta''', \beta'') < \sigma$, we may apply Proposition 5 to conclude that $[\beta'']_\varepsilon = [\beta''']_\varepsilon = [\beta]_\varepsilon$. Finally:

$$d([\beta']_\varepsilon, [\beta]_\varepsilon) = d([\beta']_\varepsilon, [\beta'']_\varepsilon) = \left|[\overline{\beta'} * \beta'']_\varepsilon\right| = |[z_n, y_n]_\varepsilon| = d(z_n, y_n) < \sigma$$

■

Theorem 29 *For every $\varepsilon, \sigma > 0$ there is a first degree polynomial $p(\sigma, \varepsilon)$ with non-negative coefficients such that $p \to 0$ as $\sigma \to 0$ (with ε fixed) and the following property holds. Let X, Y be geodesic spaces such that $\phi_{\varepsilon\omega_0} : Y_{\omega_0} \to Y_\varepsilon$ is an injection and $\omega_0 < \delta < \varepsilon$. If $f : X \to Y$ is a basepoint preserving σ-isometry for some positive constant $\sigma < \min\left\{\varepsilon - \delta, \frac{\delta - \omega_0}{4}\right\}$ then $f_\# : X_\delta \to Y_\varepsilon$ is a $p(\sigma, \varepsilon)$-isometry.*

Proof. Since $4\sigma < \delta - \omega_0$, we can choose ω such that $\omega_0 < \omega < \delta - 4\sigma$. Now we have the remaining two conditions, $\omega < \delta$ and $\sigma < \frac{\delta - \omega}{4}$, that are needed to apply Proposition 28. Note that since Y is geodesic and hence all maps ϕ_{ab} are surjective, $\phi_{\varepsilon\omega_0}$ is an isometry. But then $\phi_{\varepsilon\omega}$ is also an isometry. Since X is geodesic, we may always refine a δ-chain to an ω-chain of the same length, which means that on the right side of the inequality in the second part of Proposition 28, we may replace $|[\alpha]_\omega|$ by $|[\alpha]_\delta|$. Also since $\phi_{\varepsilon\omega}$ is an isometry, $|[f(\alpha)]_{\omega+\sigma}| = |[f(\alpha)]_\varepsilon| = |f_\#([\alpha]_\delta)|$. That is, we have

$$|[\alpha]_\delta| - |f_\#([\alpha]_\delta)| \leq \sigma\left(\frac{4}{\varepsilon} + \frac{16\sigma}{\varepsilon^2}\right)(|[\alpha]_\delta|) + \sigma\left(\frac{4\sigma}{\varepsilon} + 1\right).$$

The first part of Proposition 28 gives us:

$$|f_\#([\alpha]_\delta)| - |[\alpha]_\delta| \leq \frac{4\sigma}{\varepsilon}|[\alpha]_\delta| + \sigma.$$

Let $m(\sigma,\varepsilon) := \sigma\left(\frac{4}{\varepsilon} + \frac{16\sigma}{\varepsilon^2}\right) > \frac{4\sigma}{\varepsilon}$, $b(\sigma,\varepsilon) := \sigma\left(\frac{4\sigma}{\varepsilon}+1\right) > \sigma$ and $p(\sigma,\varepsilon)(t) = m(\sigma,\varepsilon)t + b(\sigma,\varepsilon)$. The coefficients of p then have the desired property. Since

$$d(f_\#([\alpha]_\delta), f_\#([\beta]_\delta)) = d([f(\alpha)]_\varepsilon, [f(\beta)]_\varepsilon) = \left|\left[\overline{f(\alpha)} * f(\beta)\right]_\varepsilon\right| = |[f(\overline{\alpha} * \beta)]_\varepsilon|$$

and $d([\alpha]_\delta, [\beta]_\delta) = |[\overline{\alpha} * \beta]_\delta|$ we see that $f_\#$ has distortion at most $p(\sigma,\varepsilon)$.

Finally, let $[\beta]_\varepsilon \in Y_\varepsilon$; since Y is geodesic (and all maps ϕ_{ab} are surjective) there is some $(\omega - 3\sigma)$-chain β' such that $[\beta]_\varepsilon = [\beta']_\varepsilon$. The proof is now finished by the third part of Proposition 28. ■

We need an equivariant version of pointed Gromov-Hausdorff convergence. Origins of such an idea may be found in [10], [11], and something like this was used, for example, in [9]. Let $f : X \to Y$ be a function between metric spaces and suppose there are groups H and K of isometries on X and Y, respectively, with a homomorphism $\psi : H \to K$. As usual, we say f is equivariant (with respect to ψ) if for all $h \in H$ and $x \in X$, $f(h(x)) = \psi(h)(f(x))$. Then there is a well-defined induced mapping $f_\pi : X/H \to Y/K$ defined by $f_\pi(Hx) = Kf(x)$, where $Hx := \{h(x) : h \in H\}$ is the orbit of x. We will take the quotient pseudo-metric on X/H and Y/K:

$$d(Hx, Hy) = \inf\{d(k(x), h(y)) : h, k \in H\} = \inf\{d(x, h(y)) : h \in H\}.$$

When the orbits of the action are closed sets, d is a *bona fide* metric.

In the next theorem note that some $\omega_0 > \varepsilon$ to satisfy the hypothesis always exists since the homotopy critical values are discrete. The requirement that $\frac{\varepsilon}{2} < \omega_0$ isn't firm, but it simplifies the calculation. All that really matters is that the statement is true for every $\delta < \varepsilon$ that is sufficiently close to ε.

Theorem 30 *Let $\{X_i\}$ be a collection of compact geodesic spaces such that for all i there is a basepoint preserving σ_i-isometry $f_i : X_i \to X$ for some sequence of constants $\sigma_i \to 0$. Let $\varepsilon > 0$, suppose that $\phi_{\varepsilon\omega_0} : X_{\omega_0} \to X_\varepsilon$ is injective and let $\frac{\varepsilon}{2} < \omega_0 < \delta < \varepsilon$. Then for all large i, the following hold.*

1. *$(f_i)_\# : (X_i)_\delta \to X_\varepsilon$ is a $p(\sigma_i,\varepsilon)$-isometry, and in particular, $(X_i)_\delta$ is Gromov-Hausdorff pointed convergent to X_ε.*

2. *The restriction of $(f_i)_\#$ to $\pi_\delta(X_i)$ - denoted hereafter by $(f_i)_\theta$ - is an isomorphism onto $\pi_\varepsilon(X)$.*

3. *$(f_i)_\#$ is equivariant with respect to $(f_i)_\theta$.*

Proof. The first part is an immediate consequence of Theorem 29. One need only observe that $\sigma_i < \min\left\{\varepsilon - \delta, \frac{\delta - \omega_0}{4}\right\}$ for all large i. For the next two parts, note that since f_i is basepoint preserving, $(f_i)_\theta$ does map into $\pi_\varepsilon(X)$. That $(f_i)_\theta$ is a homomorphism, and that $(f_i)_\#$ is equivariant, both follow from Equation (2), and it remains to be shown that $(f_i)_\theta$ is an isomorphism for large i . To do this we will add a few more conditions that are satisfied for all large i. We first require $\sigma_i < \frac{\varepsilon}{6}$. Next note that if we let $p(\sigma_i,\varepsilon)(t) := m_i t + b_i$, then $m_i, b_i \to 0$.

If $d\big((f_i)_\#([\alpha]_\delta),(f_i)_\#([\beta]_\delta)\big) = D$ then $d([\alpha]_\delta,[\beta]_\delta) \le \frac{D+b_i}{1-m_i}$. In particular, we may conclude the following for all large i:

$$\text{If } d((f_i)_\#([\alpha]_\delta),(f_i)_\#([\beta]_\delta)) < \frac{\varepsilon}{3}, \text{ then } d([\alpha]_\delta,[\beta]_\delta) < \frac{\varepsilon}{2}. \tag{3}$$

For large enough i the following also hold:

$$\text{If } [\beta]_\varepsilon \in X_\varepsilon, \text{ then there is some } [\alpha]_\delta \text{ such that } d((f_i)_\#([\alpha]_\delta),[\beta]_\varepsilon) < \frac{\varepsilon}{6}. \tag{4}$$

$$\text{If } d([\alpha]_\delta,[\beta]_\delta) < \frac{\varepsilon}{3}, \text{ then } d((f_i)_\#([\alpha]_\delta),(f_i)_\#([\beta]_\delta)) < \frac{\varepsilon}{2}. \tag{5}$$

Suppose that $[\lambda]_\delta \in \ker(f_i)_\theta$. Then $d((f_i)_\theta([\lambda]_\delta),[*]_\varepsilon) = 0 < \frac{\varepsilon}{3}$ and by (3), $d([\lambda]_\delta,[*]_\delta) < \frac{\varepsilon}{2} < \delta$. But ϕ_δ is injective on $B(*,\delta)$ and since λ is a loop, $[\lambda]_\delta = [*]_\delta$.

Let $[\lambda]_\varepsilon \in \pi_\varepsilon(X)$. By (4) there is $[\alpha]_\delta$ such that $d\big((f_i)_\#([\alpha]_\delta),[\lambda]_\varepsilon\big) < \frac{\varepsilon}{6}$. Letting $\alpha := \{* = x_0,\dots,x_n\}$ we have

$$d(x_n,*) \le d(f(x_n),*) + \sigma_i \le d((f_i)_\#([\alpha]_\delta),[\lambda]_\varepsilon) + \sigma_i < \frac{\varepsilon}{3} < \delta.$$

Then $\lambda' := \{* = x_0,\dots,x_n,*\}$ is a δ -loop with $d([\lambda']_\delta,[\alpha]_\delta) \le d(x_n,*) < \frac{\varepsilon}{3}$. Therefore by (5) $d\big((f_i)_\#([\alpha]_\delta),(f_i)_\#([\lambda']_\delta)\big) < \frac{\varepsilon}{2}$. The triangle inequality now shows $d\big([\lambda]_\varepsilon,(f_i)_\#([\lambda']_\delta)\big) < \frac{5\varepsilon}{6} < \varepsilon$. But once again, the injectivity of ϕ_ε on open ε-balls implies that $[\lambda]_\varepsilon = (f_i)_\#([\lambda']_\delta) = (f_i)_\theta([\lambda']_\delta)$, finishing the proof of surjectivity. ∎

Proposition 31 *Suppose that X and Y are metric spaces with groups H,K acting on X,Y, respectively, by isometries with closed orbits and $\phi: H \to K$ is an epimorphism. If $f: X \to Y$ is a σ-isometry for some first degree polynomial σ and equivariant with respect to ϕ, then $f_\pi: X/H \to Y/K$ is a σ-isometry.*

Proof. For any $x,y \in X$, we have $D := d(f_\pi(Hx),f_\pi(Hy)) = d(Kf(x),Kf(y))$, from which we obtain

$$\begin{aligned} D &= \inf\{d(f(x),h(f(y))) : h \in K\} \\ &= \inf\{d(f(x),\phi(g)(f(y))) : g \in H\} \\ &= \inf\{d(f(x),f(g(y))) : g \in H\} \end{aligned}$$

(The first equality follows because ϕ is surjective.) Now for any $g \in H$,

$$d(x,g(y)) - \sigma(d(x,g(y)) \le d(f(x),f(g(y))) \le d(x,g(y)) + \sigma(d(x,g(y)).$$

Letting $D' := d(Hx,Hy) = \inf\{d(x,g(y))\}$, the infimum of the right side is $D' + \sigma(D')$. For the left side, for arbitrary $\varepsilon > 0$ we may suppose that

$$D' \le d(x,g(y)) < D' + \varepsilon$$

which gives us

$$d(x, g(y)) - \sigma(d(x, g(y)) > D' - \sigma(D' + \varepsilon)$$

and therefore the infimum of the left side is $D' - \sigma(D)'$. That is, the distortion of f_π is at most σ.

Finally, for any $Ky \in Y/K$, there is some $f(x) = z \in Y$ such that $d(z, y) \leq \sigma(0)$. But then by definition $Kz = f(Kx)$ and $d(Kz, Ky) \leq d(z, y) \leq \sigma(0)$. ■

Proof of Theorems 1 and 2. Let $f_i : X_i \to X$ be basepoint-preserving σ_i-isometries with constants $\sigma_i \to 0$. Since the homotopy critical values of X are discrete, we may choose ω_0, and hence δ with $\frac{\varepsilon}{2} < \delta < \varepsilon$, so that the assumptions of Theorem 30 are satisfied. Eliminating finitely many terms if needed, we obtain the following properties for all i that we will use now and below: (1) $(f_i)_\# : (X_i)_\delta \to X_\varepsilon$ is a $p(\sigma_i, \varepsilon)$-isometry. (2) The restriction $(f_i)_\theta$ of $(f_i)_\#$ to $\pi_\delta(X_i)$ is an isomorphism onto $\pi_\varepsilon(X)$. (3) $(f_i)_\#$ is equivariant with respect to $(f_i)_\theta$. The first statement of Theorem 1 is an immediate consequence of (1) and (2). If ε is not a homotopy critical value of X then there are some $\varepsilon' > \varepsilon > \omega_0$ such that $\phi_{\varepsilon'\omega_0}$ is an isometry. We may now apply the first part of the theorem using ε' to see that $(X_i)_\varepsilon \to (X)_{\varepsilon'}$ and $\pi_\varepsilon(X_i)$ is eventually isomorphic to $\pi_{\varepsilon'}(X)$. But since $\phi_{\varepsilon'\varepsilon} : X_\varepsilon \to X_{\varepsilon'}$ is an isometry, $\theta_{\varepsilon'\varepsilon} : \pi_\varepsilon(X) \to \pi_{\varepsilon'}(X)$ is an isomorphism. This completes the proof of Theorem 1.

For the proof of Theorem 2 we will use the same notation and δ as chosen in the above paragraph. By eliminating terms if needed we may assume that $\varepsilon_i > \delta$ for all i. By Proposition 24 the covering space $X^{\mathcal{T}_i}_{\varepsilon_i}$ is isometrically equivalent to $X^{\mathcal{S}_i}_\delta$, and $\pi^{\mathcal{T}_i}_{\varepsilon_i}(X)$ is isomorphic to $\pi^{\mathcal{S}_i}_\delta(X)$, where $\mathcal{S}_i$ is obtained by adding to $\mathcal{T}_i$ one representative for each essential τ-circle with $\varepsilon_i > \tau \geq \varepsilon$. Therefore we need only show that there is some collection $\mathcal{T}$ as in the statement of the theorem, and a subsequence such that $(X_{i_k})^{\mathcal{S}_{i_k}}_{\delta_{i_k}} \to X^{\mathcal{T}}_\varepsilon$ and $\pi^{\mathcal{S}_{i_k}}_{\delta_{i_k}}(X_{i_k})$ is eventually isomorphic to $\pi^{\mathcal{T}}_\varepsilon(X)$. Let $g_i : (X_i)_\delta \to X_\delta$ denote $\phi^{-1}_{\varepsilon\delta} \circ (f_i)_\#$, which is a $p(\sigma_i, \varepsilon)$-isometry, and let h_i be the restriction of g_i to $\pi_\delta(X_i)$. By Conditions (2) and (3) above, h_i is an isomorphism onto $\pi_\delta(X)$ for all i, and the maps g_i are invariant with respect to the isomorphisms h_i. According to Theorem 11 of [15], the number of homotopy critical values $\geq \varepsilon$ (counted with multiplicity) in the Gromov-Hausdorff precompact collection $\{X_i\}$ has a uniform upper bound. Therefore by removing equivalent essential triads and taking a subsequence if necessary, we may assume that for some n, $\mathcal{S}_i = \{T_{i1}, \ldots, T_{in}\}$ (n could be 0, in which case the following statements about $\mathcal{T}_i$ are true for the empty set). Suppose that $T_{ij} = \{x^0_{ij}, x^1_{ij}, x^2_{ij}\}$ is a δ_{ij}-triad and T'_{ij} is an $\frac{\varepsilon}{3}$-refinement of $\alpha_{T_{ij}}$. Since the diameters of the spaces X_i have a uniform upper bound, the number of points needed to refine each $\alpha_{T_{ij}}$ has a uniform upper bound; by adding points if necessary we may assume that for some fixed w, $T'_{ij} = \{z^0_{ij} = x^0_{ij}, \ldots, z^w_{ij} = x^0_{ij}\}$ for all i, j. The uniform upper bound on diameters also implies that for some fixed m we may find $\frac{\varepsilon}{3}$-chains $\alpha_{ij} := \{* = y^0_{ij}, \ldots, y^m_{ij} = x^0_{ij}\}$ for all i, j (i.e. subdivide geodesics). Finally, let $\lambda_{ij} := \alpha_{ij} * T'_{ij} * \overline{\alpha_{ij}}$.

By choosing a subsequence yet again we may suppose that for all j, k, $f_i(z_{ij}^k) \to z_j^k$, $f_i(x_{ij}^k) \to x_j^k$ and $f_i(y_{ij}^k) \to y_j^k$. Let $\alpha_j := \{y_j^0, \ldots, y_j^m\}$, $T_j := \{x_j^0, x_j^1, x_j^2\}$, $T_j' := \{z_j^0, \ldots, z_j^w\}$, and $\lambda_j := \alpha_j * T_j' * \overline{\alpha_j}$. Since all the limiting chains have the property that each point is of distance at most $\frac{\varepsilon}{3}$ from its successor, Proposition 5 implies that there is some N such that if $i \geq N$, then

$$[f_i(\lambda_{ij})]_\rho = [\lambda_j]_\rho \text{ for all } j \text{ and any } \rho \geq \frac{\varepsilon}{2} > \frac{\varepsilon}{3}. \tag{6}$$

We assume $i \geq N$ in what follows. One immediate consequence of (6) is that

$$(f_i)_\theta\big([\lambda_{ij}]_\delta\big) = (f_i)_\#\big([\lambda_{ij}]_\delta\big) = [f(\lambda_{ij})]_\varepsilon = [\lambda_j]_\varepsilon \tag{7}$$

where $(f_i)_\theta$ is the restriction of $(f_i)_\#$ to $\pi_\delta(X_i)$.

We can now argue that T_j is an essential δ_j-triad where $\delta_j := \lim \delta_{ij} \geq \varepsilon$. In fact, by continuity of the distance function, $\delta_j = \lim \delta_{ij} \geq \varepsilon$ exists and T_j is a δ_j-triad. If T_j' were ε-null then $f_i(\lambda_{ij})$ would be also ε-null, which by (7) means $[\lambda_{ij}]_\delta \in \ker\, (f_i)_\theta$. Since $(f_i)_\theta$ is an isomorphism, λ_{ij}, and hence T_{ij}', is ε-null. This contradicts that T_{ij} is δ_{ij}-essential with $\delta_{ij} > \varepsilon$. So T_j' is not ε-null, and $\delta_j \geq \varepsilon$ by Corollary 8. If T_j were not essential then T_j' would be δ_j-null. Then for some $\delta' < \delta_j$ and close enough to δ_j, T' would also be δ'-null. But then for large enough i, $\delta_{ij} > \delta'$ and (6) implies T_{ij}' is δ'-null, a contradiction to the fact that T_{ij}' is not δ_{ij}-null.

The next consequence of (7), and the characterization of $K_\delta(\mathcal{T})$ in Remark 18, is that $h_i(K_\delta(\mathcal{S}_i)) = K_\delta(\mathcal{T})$, where $\mathcal{T} := \{T_1, \ldots, T_n\}$. At this point we may assume the following, having chosen subsequences several times (but avoiding double subscripts for simplicity): the functions $g_i : (X_i)_\delta \to X_\delta$ are $p(\sigma_i, \varepsilon)$-isometries and the restrictions k_i of h_i to $K_\delta(\mathcal{S}_i)$ are isometries onto $K_\delta(\mathcal{T})$ that are equivariant with respect to g_i. By Propositions 27 and 31, $(X_i)_\delta^{\mathcal{S}_i} = (X_i)_\delta / K_\delta(\mathcal{S}_i) \to X_\delta / K_\delta(\mathcal{T}) = X_\delta^{\mathcal{T}}$. By the choice of δ, $\phi_{\varepsilon\delta}$ is an isometry, so $X_\delta / K_\delta(\mathcal{T})$ is isometric to $X_\varepsilon / K_\varepsilon(\mathcal{T}) = X_\varepsilon^{\mathcal{T}}$ by Proposition 24.

Finally, recall that $h_i : \pi_\delta(X_i) \to \pi_\delta(X)$ is an isomorphism that takes $K_\delta(\mathcal{S}_i)$ to $K_\delta(\mathcal{T})$. Combining this with the first part of Theorem 22 gives us that $\pi_\delta^{\mathcal{S}_i}(X_i) = \pi_\delta(X_i)/K_\delta(\mathcal{S}_i)$ is isomorphic to $\pi_\delta(X)/K_\delta(\mathcal{T})$. But $\theta_{\varepsilon\delta}$ is an isomorphism from $\pi_\delta(X)$ to $\pi_\varepsilon(X)$ taking $K_\delta(\mathcal{T})$ to $K_\varepsilon(\mathcal{T})$, completing the proof of the theorem. ■

References

[1] Berestovskiĭ, Valera; Plaut, Conrad, Covering group theory for topological groups. Topology Appl. 114 (2001), no. 2, 141–186. MR1834930

[2] Berestovskiĭ, Valera; Plaut, Conrad, Uniform universal covers of uniform spaces. Topology Appl. 154 (2007), no. 8, 1748–1777. MR2317077

[3] Berard-Bergery, L., Laplacien et geodesiques ferme es sur les formes despace hyperbolique compactes, Seminaire Bourbaki, 24ieme Annee, 1971-72, Exp. 406.

[4] Burago, Dmitri; Burago, Yuri; Ivanov, Sergei, *A course in metric geometry.* Graduate Studies in Mathematics, 33. American Mathematical Society, Providence, RI, 2001. MR1835418

[5] Colin de Verdiere, Y, Spectre du laplacien et longueurs des geodesiques periodiques I, II, Compositio Math. 27 (1973), 83-106 and 159-184. MR0348798

[6] de Smit, B.; Gornet, R; Sutton, C., Sunada's method and the covering spectrum, J. Diff. Geom. 86 (2010) 501-537. MR2785841

[7] de Smit, B.; Gornet, R.; Sutton, C. Isospectral surfaces with distinct covering spectra via Cayley graphs. Geom. Dedicata 158 (2012), 343–352. MR2922720

[8] Duistermaat, J. J.; Guillemin, V. W., The spectrum of positive elliptic operators and periodic bicharacteristics, Invent. Math. 29 (1) (1975) 3979. MR0405514

[9] Ennis, John; Wei, Guofang. Describing the universal cover of a compact limit. Differential Geom. Appl. 24 (2006), no. 5, 554–562. MR2254056

[10] K. Fukaya, Theory of convergence for Riemannian orbifolds, Japan. J. Math. (N.S.) 12 (1) (1986) 121–160. MR0914311

[11] K. Fukaya, T. Yamaguchi, The fundamental groups of almost non-negatively curved manifolds, Ann. of Math. (2) 136 (2) (1992) 253–333.

[12] Munkres, J., *Topology (2nd. Ed.)*, Prentice Hall, 2000. MR1185120

[13] Plaut, Conrad Metric spaces of curvature $\geq k$. Handbook of geometric topology, 819–898, North-Holland, Amsterdam, 2002. MR1886682

[14] Plaut, Conrad, Quotients of uniform spaces. Topology Appl. 153 (2006), no. 14, 2430–2444. MR2243722

[15] Plaut, Conrad; Wilkins, Jay, Discrete homotopies and the fundamental group. Adv. Math. 232 (2013), no. 1, 271–294. MR2989983

[16] Sormani, Christina; Wei, Guofang, Hausdorff convergence and universal covers. Trans. Amer. Math. Soc. 353 (2001), no. 9, 3585–3602. MR1837249

[17] Sormani, Christina; Wei, Guofang, The covering spectrum of a compact length space. J. Differential Geom. 67 (2004), no. 1, 35–77. MR2153481

[18] Sormani, Christina; Wei, Guofang, Universal covers for Hausdorff limits of noncompact spaces. Trans. Amer. Math. Soc. 356 (2004), no. 3, 1233–1270. MR2021619

[19] Sormani, Christina; Wei, Guofang. The cut-off covering spectrum. Trans. Amer. Math. Soc. 362 (2010), no. 5, 2339–2391. MR2584603

[20] Spanier, Edwin H., Algebraic topology. Corrected reprint. Springer-Verlag, New York-Berlin, 1981. xvi+528 pp. ISBN: 0-387-90646-0 MR0666554

[21] Sunada, T., Riemannian coverings and isospectral manifolds, Ann. of Math. (2) 121 (1985), no. 1, 169-186. MR0782558

[22] Wilkins, Jay, The revised and uniform fundamental groups and universal covers of geodesic spaces. Topology Appl. 160 (2013), no. 6, 812–835. MR3028614